


Technical Report
1084

Two-Dimensional ESPRIT with Tracking for Radar Imaging and Feature Extraction

M. L. Burrows

2 August 2002

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



Prepared for the Department of the Army under Air Force
Contract F19628-00-C-0002.

Approved for public release; distribution is unlimited.

20020816 071

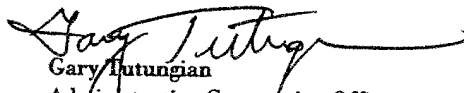
This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. The work was sponsored by the U.S. Army Space & Missile Defense Command under Air Force Contract F19628-00-C-0002.

This report may be reproduced to satisfy needs of U.S. Government agencies.

The ESC Public Affairs Office has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER


Gary Tutungian
Administrative Contracting Officer
Plans and Programs Directorate
Contracted Support Management

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission is given to destroy this document
when it is no longer needed.

Massachusetts Institute of Technology
Lincoln Laboratory

**Two-Dimensional ESPRIT with Tracking for
Radar Imaging and Feature Extraction**

M.L. Burrows
Group 34

Technical Report 1084

2 August 2002

Approved for public release; distribution is unlimited.

ABSTRACT

ESPRIT processing appears to be the best of the known spectral-analysis techniques. It provides the highest resolution and has no spectral splatter. By applying matrix eigenstructure analysis, it gives a direct answer to the direct question "What frequencies, real or complex, are present in the data and what are their amplitudes?" Conventional Fourier techniques, as well as some of the other higher-resolution methods, answer the less direct question "What amplitudes, applied to a set of regularly-spaced real frequencies, best represent the data?" Then comes the problem of interpreting the amplitudes.

These attributes of ESPRIT, in the two-dimensional version described here, make it a natural for radar signal processing, where it answers the need for high-resolution imaging, free of sidelobes in range and Doppler, and for high-fidelity target feature extraction. For example, the uncertainty in the scatterer-center locations in an ESPRIT image extracted from high-quality static-range radar data collected over a bandwidth of 1 GHz is just a few millimeters; for conventional Fourier processing of the same data the uncertainty is many centimeters. The signature of the base edge of a perfectly conducting cone extracted from static-range data by ESPRIT agrees accurately with the signature predicted by edge-diffraction theory.

This report starts with a mathematical model for the radar data, describes a technique for "resampling" the data to achieve a more perfect fit with the ESPRIT data model, summarizes the two-dimensional ESPRIT algorithm itself, and presents examples of its performance. The appendix covers the details of this least-mean-square version of ESPRIT, including an enhancement that allows the scatterers to be tracked individually. The algorithm properly distinguishes between target locations having one coordinate in common, and it automatically associates correctly in pairs each entry in the list of ranges with the corresponding entry in the list of range rates.

ACKNOWLEDGMENTS

I would like to thank Joe Mayhan for his continuing inspiration and support, Henry Helmken for providing an accurate mother lode of static-range data, and Jean Piou for alerting me to the early emergence of ESPRIT-like ideas in the systems theory literature.

TABLE OF CONTENTS

	Page
Abstract	iii
Acknowledgments	v
List of Illustrations	ix
1. INTRODUCTION	1
2. DATA MODEL AND RESAMPLING	5
3. TWO-DIMENSIONAL ESPRIT	7
4. MORE RESULTS	9
5. CONCLUSIONS	15
APPENDIX. THE ESPRIT DETAILS	17
REFERENCES	23

LIST OF ILLUSTRATIONS

Figure No.		Page
1	The Fourier and ESPRIT X-band radar images of an aluminum round-nosed cone target at nose-on aspect derived from measured static-range data of 1 GHz bandwidth. The target's rotation axis and the radar's magnetic-field polarization on transmit and receive were vertical (HH polarization); the target's body-axis of symmetry and the radar line of sight were horizontal.	2
2	The complex diffraction coefficient of the base edge extracted by ESPRIT processing, from the same measured static-range HH-polarization data used for Figure 1, compared with the coefficient predicted by the geometric theory of diffraction.	3
3	ESPRIT X-band images from different aspects and 360-composite images, both in B&W and, to include amplitude information, in false color. The images were extracted from the same measured static-range X-band HH-polarization radar data used for the previous figures.	9
4	Narrow-angle imaging with ESPRIT, including zoom views showing the smaller location uncertainty that greater signal-to-noise ratio provides.	11
5	360° trajectories of the individual tracked scatterers, in slant range and cross range. Apparent is the narrowing, in the vicinity of nose-on incidence and "base-on" incidence, of the cross-range spacing of the pairs of scatterers from the base edge and from slot 3, consistent with Figure 1. The bottom panel shows the corresponding backscattering cross sections of the individual scatterers. The top one shows the normalized mean square error between the array of points in the static-range processing window and the array of points reconstructed from the point-scatterer model of the target implied by the ESPRIT analysis.	12
6	The backscattering cross section of the target as a function of aspect, at the center frequency of the band covered by the data, compared with its backscattering cross section reconstructed from the ESPRIT analysis.	13
7	Scatterer trajectories for aspects within $\pm 30^\circ$ of nose-on aspect.	14

1. INTRODUCTION

Mathematically, a range-Doppler radar image is created by applying two-dimensional spectral analysis to a sequence of wideband radar returns. The two dimensions are time and frequency; they express the variation in time over the whole sequence of returns of the amplitude and phase of the frequency components in each return separately. Generally speaking, each two-dimensional spectral component identified in the sequence of returns is associated with a particular scattering center on the target. Its phase behavior, as a function of frequency and time, determines the slant range and cross range of the scattering center. Its amplitude coefficient determines the effective radar cross section of the scatterer.

This spectral analysis is conventionally performed today by carrying out a two-dimensional digital Fourier transform on a rectangular array of samples taken from the radar returns. Over the last decade or so, the more powerful spectral analysis technique known as ESPRIT (Estimation of Signal Parameters via Rotational Invariance Techniques) [1,2] has been developed that achieves a very significant improvement in spectral resolution. Other important advantages of ESPRIT are that there are no range or Doppler sidelobes to confuse image interpretation, and that ESPRIT can be used to extract accurately the radar signatures, in amplitude and phase, of the individual scattering features of the radar target.

Both the Fourier transform and ESPRIT methods of spectral analysis recast the data into a weighted sum of complex exponentials of different frequencies. The advantages of ESPRIT arise from its singular ability both to identify and quantify with precision the actual frequency components present in the data. It estimates the complex amplitude and frequency of each component, accepting even frequencies that are complex, a feature that allows decaying or growing exponentials to be accommodated. The Fourier transform, in contrast, achieves its Procrustian fit to the data by adjusting only the amplitudes of a preset list of regularly spaced real frequencies.

Figure 1 compares, for an aluminum cone target, the X-band images extracted by ESPRIT and by Fourier processing from the same block of static-range backscattering data. The data block was selected to cover only a narrow angular sector centered on nose-on aspect; the bandwidth of the data was 1 GHz. The ESPRIT image consists simply of the estimated locations of the principal sources of the scattered field. The Fourier image, since these locations are not directly available, is the conventional contour plot of the discrete two-dimensional Fourier transform of the windowed radar data. The improvement that ESPRIT processing provides, both in sharpness and spectral splatter, is striking. It sees clearly and places accurately all three slots, resolves reliably the two "edges" of slot 3, the largest diameter one, and recognizes, though not accurately delineating, the extended nature of slot 2. (The phantom scatterer lying behind the base of the cone is the result of the doubly diffracted ray that crosses the flat base of the target from one edge to the other [3].)

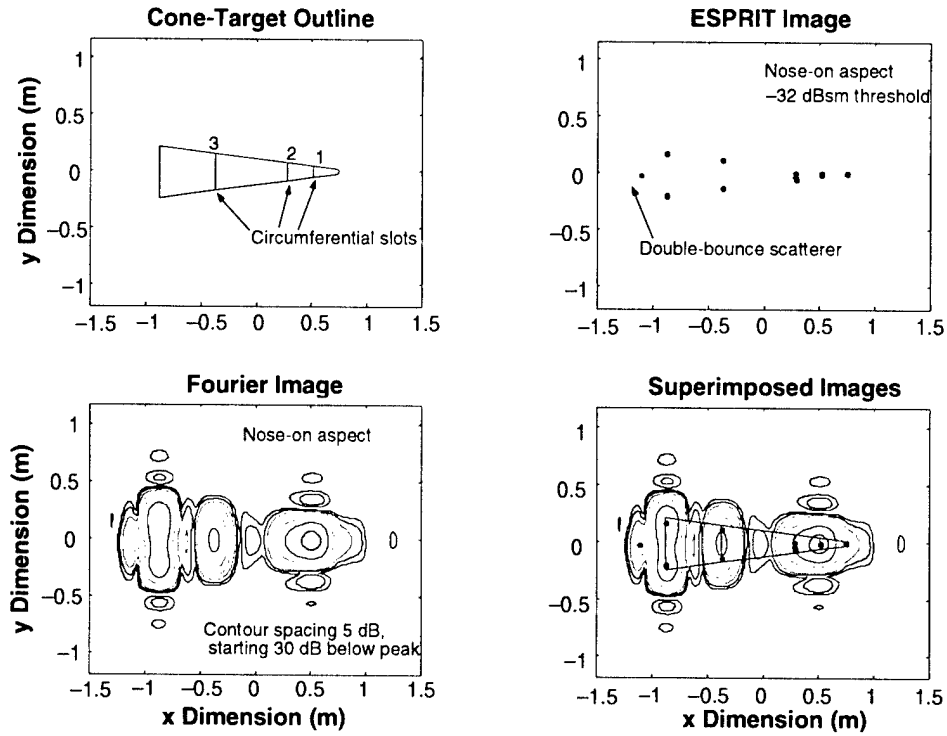


Figure 1. The Fourier and ESPRIT X-band radar images of an aluminum round-nosed cone target at nose-on aspect derived from measured static-range data of 1 GHz bandwidth. The target's rotation axis and the radar's magnetic-field polarization on transmit and receive were vertical (HH polarization); the target's body-axis of symmetry and the radar line of sight were horizontal.

Figure 2 shows, for the full 360° extent of the same data collection, the close agreement between the complex diffraction coefficient of the base edge of the target, extracted by means of ESPRIT processing, and the coefficient predicted by the geometric theory of diffraction [4]. The magnitude of the diffraction coefficient was equated directly with the absolute magnitude of the amplitude coefficient associated with the scatterer at the base edge. Its phase was evaluated by recentering, by a process of trial and error, the phase of the amplitude coefficient from the axis of target rotation to the base edge of the target. Since the target motion was that of pure rotation, this process involved simply multiplying the complex amplitude coefficient of the base edge by the factor $\exp\{i2kr \cos(\theta - \theta_0)\}$, in which r and θ_0 were adjusted to attain as closely as possible a piece-wise constant-phase behavior.




Figure 2. The complex diffraction coefficient of the base edge extracted by ESPRIT processing, from the same measured static-range HH-polarization data used for Figure 1, compared with the coefficient predicted by the geometric theory of diffraction.

The concept behind the ESPRIT technique appeared earlier in the systems analysis literature (see [5], for example) and was presented in the vocabulary specific to that specialty. A recent manifestation of that approach, focused more directly on the radar imaging problem, appears in [6]. The ESPRIT rendition has achieved wider recognition because it is presented in the more familiar terms of applied mathematics and its application is more direct. It was used to produce radar imaging and feature extraction results reported in [7], but was not further described there. The version presented here is an improved version of that earlier one.

There are other high-resolution spectral analysis techniques. Notable are the minimum variance method of Capon, Burg's maximum entropy method, Pisarenko's harmonic decomposition method, the Yule-Walker auto-regressive method, and Schmidt's MUSIC (Multiple Signal Classification) method. They are described and evaluated by Proakis *et al.* [8], who find that only MUSIC can match the resolution of the ESPRIT algorithm. Moreover, only ESPRIT itself estimates directly the frequencies, real or complex, present in the data. All the others depend on interpreting what is essentially a pictorial

presentation of the spectrum. Comparison can be made with the Fourier method, for example, for which the estimated amplitude coefficients define the Fourier spectrum that must then be the subject of a numerical search to evaluate the locations and amplitudes of the spectral components. This is inconvenient in one dimension; in two it is a real problem.

Some methods specifically addressing the radar resolution problem have been proposed. They are intuitively appealing, but are also circuitous in the way they seek to answer the basic question "What frequencies, real or complex, are present?" One is "bandwidth extrapolation," a data-extension method initially proposed by Bowling and successfully implemented by Cuomo [9]. In effect, it applies conventional Fourier processing to data that has been extrapolated outside the collection bandwidth. The spectral splatter problem of Fourier processing remains. Another is "high-definition vector imaging" by Benitz [10], which presents the spectral analysis problem as that of probing a field of view with a narrow antenna beam. The method attempts to overcome the spectral splatter problem by applying an adaptive nulling technique to suppress targets in the sidelobes of the beam. Again, it is a roundabout way of answering the basic question.

At first glance, the greatly improved resolution attainable with the ESPRIT technique would appear to be inconsistent with Shannon's sampling theorem, which imposes a strict limit on the number of degrees of freedom in a time- and bandwidth-limited signal. Since ESPRIT allows scatterers very close together to be distinguished from one another, could it not therefore separately resolve an arbitrarily large number of scatterers? The answer is, simply, no. Like Fourier imaging, the maximum number of scatterers it can resolve is limited by Shannon's sampling theorem. The advantage of ESPRIT is that it can accommodate a much less uniform spacing of the scatterers and still resolve them all.

The following sections start by introducing the radar data model, including a technique of "resampling" for overcoming its imperfect fit with the data model required by ESPRIT, then describing the principle of the two-dimensional ESPRIT method itself, and finishing with some more examples of the algorithm's performance. The appendix covers the details of this least-mean-square version of ESPRIT, including an enhancement that allows the individual scatterers to be tracked. The algorithm properly distinguishes between target locations having one coordinate in common, and it automatically associates correctly in pairs each entry in the lists it evaluates of ranges and range rates.

2. DATA MODEL AND RESAMPLING

The sampled backscattered radar return y_{mn} at frequency $f_0 + n\Delta f$ from a cluster of uniformly moving point targets at range $r_s + u_s m\Delta t$ can be expressed as the weighted sum $y_{mn} = \sum_s w_s \exp[i4\pi(f_0 + n\Delta f)(r_s + u_s m\Delta t) / c]$, where r_s and u_s are the range and range rate at time zero of point target s . In practice, these samples could be generated from a sequence of wideband uncompressed chirp radar returns, properly equalized in phase and amplitude to remove any nonuniformity in frequency and time of the transmitted signal and of the transfer function of the complete signal path from transmitter to received video. This model can accommodate mutual coupling between the point targets; a double reflection, for example, would cause a spurious target to appear—a target in the post-processing image of the cluster corresponding to no actual target. It can also accommodate a simple exponential variation in time or frequency in the strength of the return, which allows it to model cross-section variations in frequency and aspect. These variations would manifest themselves as non-zero imaginary parts of the range r_s and range rate u_s .

In this weighted sum formula for the data, the presence of the cross product $i4\pi mn\Delta f\Delta t u_s / c$ in the exponent prevents it from conforming strictly with the signal structure $\sum_s a_s \exp(im\alpha_s + in\beta_s)$ required for ESPRIT (or, for that matter, for conventional Fourier) spectral analysis. This term is negligible, if the time and frequency spread of the data are small enough. But tests have shown that it can have a significant negative effect on ESPRIT processing for practical values of these spreads. Fortunately, in cases where it is a problem, it can be removed by resampling the data. If the data samples are arranged in a rectangular array with m as the row number and n as the column number, this technique involves interpolating along each column of data to achieve in effect a frequency dependent sampling interval for each column. It amounts to choosing the new time sampling interval $\Delta t'$ to satisfy the constraint $(f_0 + n\Delta f)\Delta t' = f_0\Delta t$. A standard spline interpolation works well.

At frequencies below the center frequency of the data window, the resampling process increases the total time width of the data window. This implies that, to avoid asking the spline interpolation to *extrapolate* at the ends of the window, the number of samples provided initially must be greater than the number required for data processing by at least the factor $1 / [1 - B / (2f_0)]$, where B is the bandwidth of the data window and f_0 its center frequency. To maintain symmetry, the difference between the two numbers must also be even.

Resampling casts the data into the required form for the data array Y :

$$[Y]_{mn} = y_{mn} = \sum_s a_s \exp(im\alpha_s + in\beta_s), \quad (1)$$

with $\alpha_s = 4\pi f_0 u_s \Delta t / c$, $\beta_s = 4\pi r_s \Delta f / c$, and $a_s = w_s \exp(i4\pi f_0 r_s / c)$, where now Δt is the original time sampling interval and the ranges of both m and n over the rows and columns of Y are centered on

zero. The ESPRIT algorithm estimates the parameters $\exp(i\alpha_s)$ and $\exp(i\beta_s)$ for each point target s contributing to the data array by manipulating the data array Y in such a way that they become the solutions of two coupled eigenvalue problems. From them, the extraction of the range r_s and range-rate u_s is a straightforward application of the above identities. The addition of a motion model for the target then determines the cross range, and the complex amplitude a_s is evaluated by a least-squares fit procedure applied to (1), the α_s and β_s now being known quantities.

Although this data model can accommodate exactly a complex of point targets moving uniformly, with or without cross coupling, and with scattering amplitudes having exponential frequency and time dependencies, it is not exact for the rotary motion common to target imaging and feature extraction. Without further processing, for these applications the data model is an approximation. On the other hand, in most cases, limiting the angle swept out by the arc of motion covered by the data set reduces the problem to insignificance.

3. TWO-DIMENSIONAL ESPRIT

Two-dimensional ESPRIT spectral analysis exploits the fact that, according to the mathematical model (1) of the data, the contributions of spectral component s to any two equally sized subarrays of the data array Y differ only by a shift factor determined by the number of rows and columns by which the subarrays are offset from one another. For an offset of p rows and q columns, that shift factor is $\exp(ip\alpha_s + iq\beta_s)$, and since it depends on s , the two subarrays are, in total, two different linear combinations of all the individual spectral contributions. Of special significance are the values $\exp(i\alpha_s)$ and $\exp(i\beta_s)$ taken by the shift factor when one of the pair p and q is one and the other zero.

This key feature of Y makes it possible to rearrange and select its elements, in the manner described in the Appendix, to define three "shift" matrices U_0 , U_l , and U_f of block-Hankel structure with the following properties:

1. Within each matrix, each column is a different linear combination of all the individual spectral contributions.
2. The rank of each matrix, in the absence of noise, is equal to the total number S of spectral components in the data. (This assumes that the smaller dimension of each matrix is no less than S .)
3. The contribution of spectral component s to U_l or U_f differs only by the shift factor $\exp(i\alpha_s)$ or $\exp(i\beta_s)$ from the contribution of the same spectral component to U_0 .

These properties imply that in the absence of noise there exists a vector x_s that can recombine the columns of U_0 , U_l , and U_f to extract the particular spectral component s and also that the result of the extraction on the three matrices differs only by the factors $\exp(i\alpha_s)$ and $\exp(i\beta_s)$. Specifically, they imply $U_l x_s = U_0 x_s \exp(i\alpha_s)$ and $U_f x_s = U_0 x_s \exp(i\beta_s)$, two coupled generalized eigenvalue problems. The exponents of their eigenvalue solutions, according to the derivation in the last section, are proportional to the range and cross range of target s .

The details of the solution are covered in the Appendix, including the initial singular-value decomposition processing of the data array to strip away noise; evaluation of the three shift matrices U_0 , U_l , and U_f ; a procedure for achieving the correct pairing of the S eigenvalues of each of the two eigenvalue problems; and a final gradient-following coupled least-mean-square fine tuning of the paired solutions, which in the presence of noise are not independently determined. This last process provides the opportunity for following the movement of the separate scatterers as the target rotates by automatically keeping track of them.

This completes the evaluation of the range/range-rate pairs for all the targets, the more computationally intensive part of the spectral analysis. The second part is the evaluation of the amplitudes

a_s of the spectral components. These are evaluated readily using standard least-mean-square techniques to fit the formula (1) for the elements of the data array to the measured data array, using the now-known values of the α_s and β_s .

4. MORE RESULTS

Some more ESPRIT X-band images from the same data are shown in Figure 3.

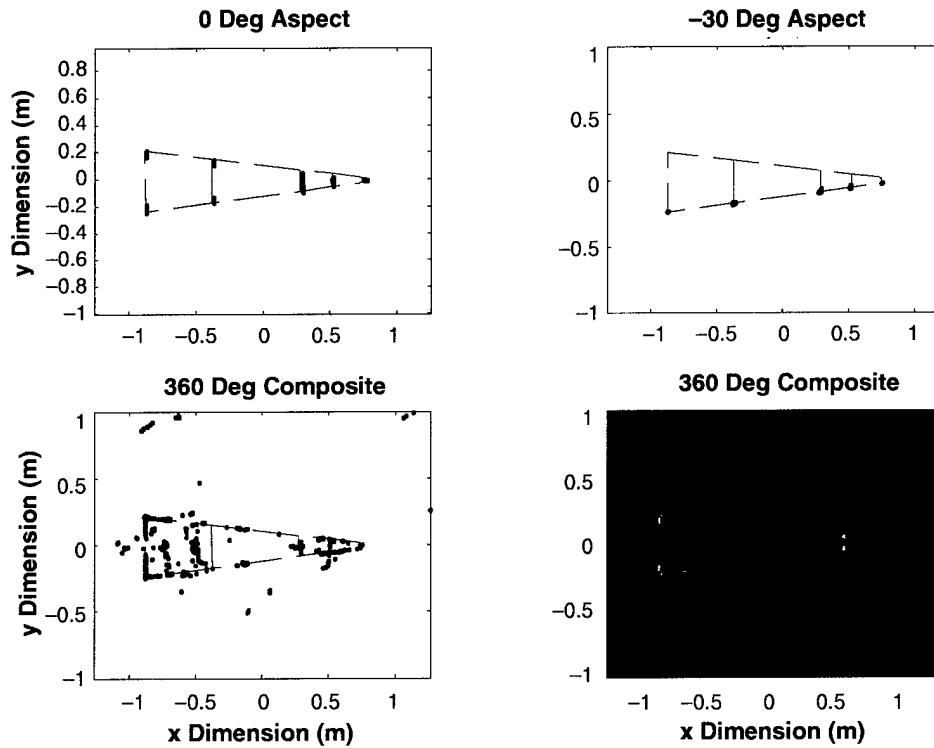


Figure 3. ESPRIT X-band images from different aspects and 360-composite images, both in B&W and, to include amplitude information, in false color. The images were extracted from the same measured static-range X-band HH-polarization radar data used for the previous figures.

The 0° aspect image (top left panel) has a higher amplitude threshold than that of Figure 1, and so the double-bounce “phantom” scatterer behind the base no longer appears. Also, a wider range of aspect angles was used, so the scatterers lie closer to the target outline. (The radar return appears to come from what is essentially a Fresnel zone in the shape of an arc of the base edge or slot. Close to nose-on aspect, this zone covers a greater length of arc, and so has an effective center of radiation displaced significantly towards the center of curvature of the arc.)

At -30° aspect (top right panel) only the scatterers along one side of the target can be seen and they are now located right at the outline of the target. Those on the other side are in shadow.

The lower two panels show the composite images that result when the images of all aspects from -180° to $+180^\circ$ are superimposed on the same plot. For the black and white image, the threshold was set high enough to reject the scatterers on opposite sides of slot 2, but the greater cross section of the base edge comes through strongly, clearly outlining its physical shape. The color composite, with the extra degree of freedom color provides, shows all the significant scatterers, including, faintly, the "phantom" scatterer seen in Figure 1.

As an illustration of the effect of signal-to-noise ratio on location uncertainty, Figure 4's three zoom views of the -30° image in Figure 3, all to the same scale, show the scatter of points generated by the ESPRIT algorithm for three different scatterers. (The processing window was moved, a fraction of a degree at a time, over a narrow range of aspects centered on the 30° aspect. At each position, a fresh estimate of the location of all the scatterers was made. Thus, there are as many points plotted for each scatterer as there were distinct processing-window positions.) The higher signal-to-noise ratio of the radar return from the base reduces the uncertainty in its estimated location. Also apparent is the smaller uncertainty in slant range than in cross range.

Since the ESPRIT algorithm, in the form described here, allows the individual scattering centers to be individually tracked, their trajectories in slant range and cross range can be plotted as separate continuous curves. Figure 5 shows the result. Each scatterer has uniquely associated with it a slant range, a cross range, and a complex diffraction coefficient over a wide extent of aspects for which it is not shadowed or too small to be seen, and so can be assigned a unique color. The measured curve in Figure 2 for the backscattering cross section of the base edge could in fact have been directly copied from Figure 5 in just two operations, one for the section extending in aspect from -190° to about -90° , and the other for the section -10° to $+190^\circ$. The cross-range plot shows the narrowing, in the vicinity of nose-on incidence and "base-on" incidence, of the spacing between the members of the pairs of scatterers associated with the base edge and with slot 3, consistent with Figure 1.

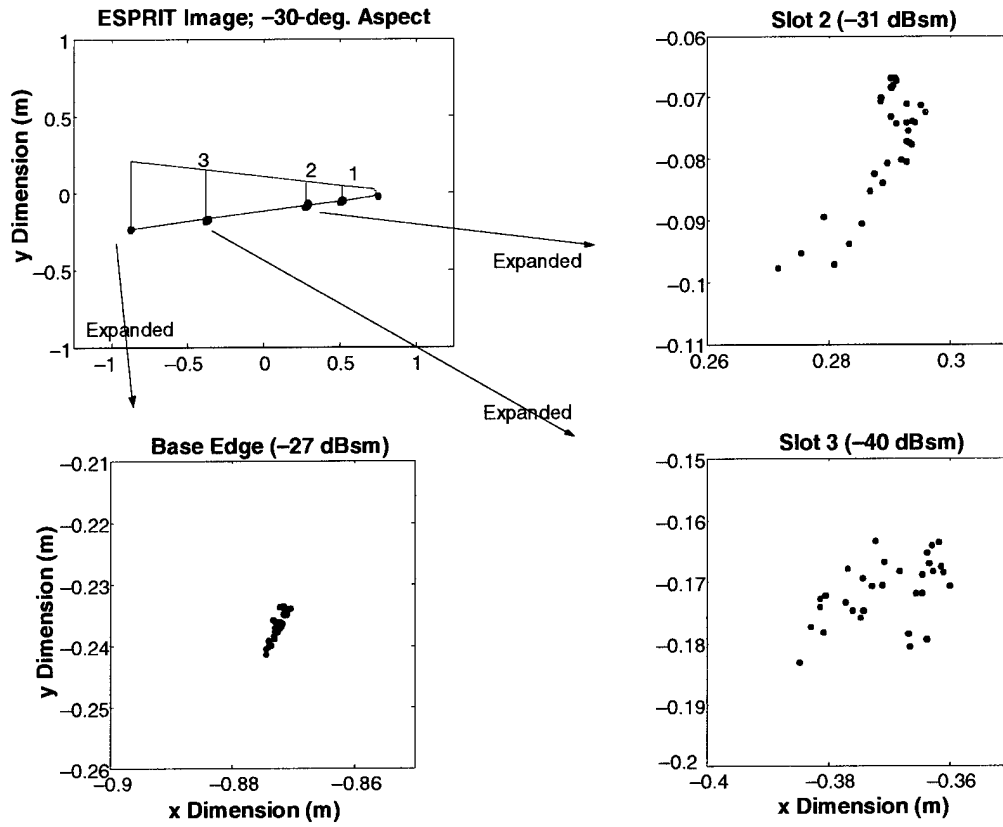


Figure 4. Narrow-angle imaging with ESPRIT, including zoom views showing the smaller location uncertainty that greater signal-to-noise ratio provides.

The upper panel in Figure 5 compares, as a function of aspect, the normalized mean square error between the array of complex data samples in the static-range processing window and the corresponding array of complex data samples reconstructed from the point-scatterer model of the target implied by the ESPRIT analysis. The error is low enough to suggest that most of the energy in the original data is accounted for. Significant is the fact that the error is low near 0° aspect and near 180° aspect, corresponding to regions in which the signal-to-noise ratio is higher and the scatterers are tracked more faithfully.




Figure 5. 360° trajectories of the individual tracked scatterers, in slant range and cross range. Apparent is the narrowing, in the vicinity of nose-on incidence and "base-on" incidence, of the cross-range spacing of the pairs of scatterers from the base edge and from slot 3, consistent with Figure 1. The bottom panel shows the corresponding backscattering cross sections of the individual scatterers. The top one shows the normalized mean square error between the array of points in the static-range processing window and the array of points reconstructed from the point-scatterer model of the target implied by the ESPRIT analysis.

Figure 6 shows an alternative way of looking at the errors in the ESPRIT analysis. It shows the backscattering cross section of the target as a function of aspect, at the center frequency (10 GHz) of the 1 GHz band covered by the data, compared with its backscattering cross section reconstructed from the ESPRIT analysis. The agreement between the two is seen to be very good. The same reassuring result is obtained when the comparison is made at any other frequency in the 1 GHz data band.




Figure 6. The backscattering cross section of the target as a function of aspect, at the center frequency of the band covered by the data, compared with its backscattering cross section reconstructed from the ESPRIT analysis.

Figure 7 is a closer view of the scatterer trajectories. It plots the same quantities as Figure 5, but restricts the aspect extent to within $\pm 30^\circ$ of nose-on. More readily visible here than in Figure 5 is the narrowing near nose-on aspect, described above, of the cross-range spacing of the pairs of scatterers associated with the base edge and with slot 3. It also shows more clearly the separate tracking of each scatterer in the pair right through nose-on aspect. The pairs of scatterers associated with slots 1 and 2 are evidently too close together to be separately resolved, and so only one scatterer is tracked for each slot, but the cross-range trajectory shows evidence of mutual interference between the two members of each pair.

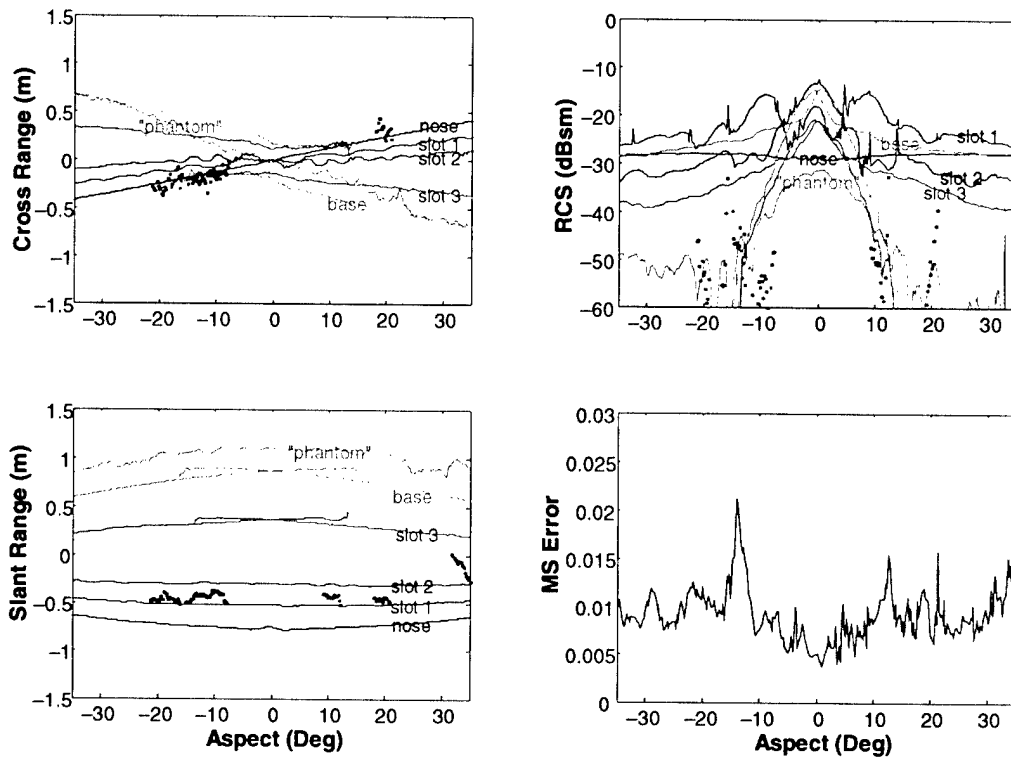


Figure 7. Scatterer trajectories for aspects within $\pm 30^\circ$ of nose-on aspect.

5. CONCLUSIONS

ESPRIT processing appears to be the best of the known spectral-analysis techniques. It provides the highest resolution and has no range or Doppler sidelobes. By applying matrix eigenstructure analysis, it gives a direct answer to the question "What frequencies, real or complex, are present in the data and what are their amplitudes?" Conventional Fourier techniques, as well as some of the other higher-resolution methods, answer the less direct question "What amplitudes, applied to a set of regularly-spaced real frequencies, best represent the data?"

These attributes of ESPRIT, in the two-dimensional version described here, make it a natural for radar signal processing, where it answers the need for high-resolution imaging and target feature extraction. It can provide a tracking capability useful both for the construction of the motion model of the target and for isolating and measuring automatically the scattering parameters of a chosen scatterer.

APPENDIX

THE ESPRIT DETAILS

THE COUPLED EIGENVALUE PROBLEMS

The input to the two-dimensional ESPRIT algorithm is the resampled data array Y having, in the absence of noise, the mathematical form defined in (1). Provided aliasing in both dimensions is avoided, this structure implies that every two-dimensional subarray of Y of a given size will be a different linear combination of the same S two-dimensional spectral components of the form $\exp(im\alpha_s + in\beta_s)$. Each spectral component arises from a different point scatterer. It follows that if the elements of each of these subarrays are rearranged in some systematic manner into a single column vector, and these column vectors are then concatenated to form a matrix H , each column of H will be a different linear combination of the same S spectral components. In addition, it follows that no two spectral components can have the same vector contribution to any column unless their corresponding point targets lie in the same place. This implies that there will exist a set of vectors x_s , one for each point scatterer, such that the column vector Hx_s will differ by only a scalar factor from the vector contribution of point scatterer s to any one of the columns of H .

The now-standard way of building H , the first step of the algorithm, is rearranging the elements of Y into block-Hankel form. This is accomplished in two steps. First the Hankel matrix form of each column of Y is created, in which element (p, q) of the Hankel matrix is element $(p + q - 1)$ of the corresponding column. Then these Hankel matrices are arrayed to form the block-Hankel matrix, in which block (p, q) of the block-Hankel matrix is the Hankel matrix formed from column $(p + q - 1)$ of Y .

The second step of the algorithm is applying singular value decomposition to reduce the second dimension of H to match its essential rank. That is, H is decomposed as $H = U\Sigma V'$ and only the columns of U corresponding to the significant singular values (the elements of the diagonal matrix Σ) are retained. (Here, following the convention used by MATLAB, the notation A' is used to denote the conjugate transpose of A .) One suitable test for significance is the size of the singular value relative to the maximum singular value. Another is the absolute size of the singular value. In the noise free case, only S non-zero singular values will exist, and so only S columns of U will be retained. Typically, the number of rows will be larger than S .

The matrix U retains the important property of H that there will exist a vector x_s such that the column vector Ux_s will differ by only a scalar factor from the vector contribution of point scatterer s to any one of the columns of H .

The third step exploits the fact that, by virtue of the block-Hankel rearrangement of Y , the contribution of point scatterer s to any one row of U differs from its contribution to the previous row in

the same block by the factor $\exp(i\alpha_s)$, and from its contribution to the corresponding row in the previous block by the factor $\exp(i\beta_s)$. This implies that by deleting specific sets of rows¹ of U , three versions U_0 , U_t , U_f of U can be created such that, in the absence of noise, $U_t x_s = \exp(i\alpha_s) U_0 x_s$ and $U_f x_s = \exp(i\beta_s) U_0 x_s$. By multiplying both sides of both equations by $(U_0' U_0)^{-1} U_0'$, these can be rewritten as the two coupled eigenvalue problems

$$(R_t - \tau_s I) x_s = 0 \quad (2a)$$

$$(R_f - \nu_s I) x_s = 0 \quad (2b)$$

where the square $S \times S$ matrices R_t and R_f are given by

$$R_h = (U_0' U_0)^{-1} U_0' U_h, \quad (h = t, f) \quad (3)$$

$\tau_s = \exp(i\alpha_s)$, $\nu_s = \exp(i\beta_s)$ and I is the identity matrix.

The eigenvalues, $\exp(i\alpha_s)$ and $\exp(i\beta_s)$, are the crucial parameters, since from them the range r_s and range-rate u_s of the scatterers are evaluated using $\alpha_s = 4\pi f_0 u_s \Delta t / c$ and $\beta_s = 4\pi r_s \Delta f / c$.

EIGENVALUE PAIRING

Solving the eigenvalue problems (2a) and (2b) separately would deliver two lists of eigenvalues. Unfortunately, the elements in the two lists would not necessarily be in the same order. This means that the pairing information, necessary for associating the eigenvalues of the two lists in pairs, one pair per scatterer, is missing. One simple way of avoiding this problem is to note that since the eigenvectors of the two eigenvalue problems are the same, the eigenvectors can be evaluated just once, as the eigenvectors of the weighted matrix sum $c_1 R_t + c_2 R_f$, where the scalar weights c_1 and c_2 can be chosen essentially arbitrarily. Then the properly paired eigenvalues corresponding to each (normalized) eigenvector x_s are given by $\tau_s = x_s' R_t x_s$ and $\nu_s = x_s' R_f x_s$ [11]. One simple choice for the weights could be $c_1 = 1$ and $c_2 = 0$, but in the presence of noise a more accurate estimate of the eigenvalues is obtained if both matrices are brought to bear by making both weights non-zero.

¹ Deleting from U the last row of each block of rows and also the whole last block of rows creates U_0 . Deleting the first row of each block and the whole last block of rows creates U_t . Deleting the last row of each block and the whole first block of rows creates U_f .

This pairing technique fails in the rare situation in which the weighted eigenvalue sum $c_1 \tau_i + c_2 \nu_j$ for one scatterer happens to be the same as that for another. If that happens, it can be avoided simply by changing the scalar weights in the weighted matrix sum $c_1 R_i + c_2 R_j$.

Another pairing technique is described by Piou *et al.* [6].

EFFECT OF NOISE

In the presence of additive noise, each right side of the two eigenvalue problems (2a) and (2b) will be an error vector rather than zero. Now the problem becomes one of finding the eigenvalue pair (λ_s, ν_s) and corresponding eigenvector x_s that together minimize the sum of the squared 2-norms of the two error vectors: $\|(R_i - \tau_s I)x_s\|_2^2 + \|(R_f - \nu_s I)x_s\|_2^2$.

The calculus of variations applied to this mean-squares minimization problem leads to the following set of simultaneous equations for the x_s , τ_s , and ν_s :

$$\left[(R_i' - \tau_s^* I)(R_i - \tau_s I) + (R_f' - \nu_s^* I)(R_f - \nu_s I) \right] x_s = 0 \quad (4a)$$

$$\tau_s = x_s' R_i x_s \quad (4b)$$

$$\nu_s = x_s' R_f x_s, \quad (4c)$$

where x_s is construed to be of unit norm.

Their solution can be found using a gradient method starting from initial estimates of the τ_s and ν_s . These initial estimates are the eigenvalues of the coupled eigenvalue problems (2a) and (2b), obtained in the manner described. With them, the matrix in square brackets in (4a) is evaluated and then factored by singular value decomposition to determine the normalized vector x_s that minimizes the norm of the left side of (4a). (This vector is equal to the post-multiplying vector, in the singular value decomposition, corresponding to the minimum singular value.) Substituting this vector in (4b) and (4c) yields the updated values of τ_s and ν_s that comprise the estimates for the next iteration. It is found in practice that the process converges quickly, making it necessary to use only a few iterations. The inclusion of this gradient method as a final step in the algorithm has shown in tests to improve significantly the precision of the final values of the τ_s and ν_s .

TRACKING

This technique of gradient convergence to a solution opens the way to a crude but effective method of tracking the individual scatterers in range and range rate as they move with the rotating target.

A direct application of the algorithm to following the motion of the scatterers would involve simply incrementing the data window in time and then, at each location, reevaluating their range and range rate. This would not constitute tracking, because a new version of the pairing problem would arise. Specifically, unanswered would be the question of which scatterer in the list of coordinates produced at the current window location corresponds to any specific scatterer in the previous list.

The procedure for overcoming this is to use the values of τ_s and ν_s for scatterer s evaluated in the previous window location as the initial estimates of these parameters for the current window location. Applying the gradient method, with values for the matrices R_i and R_f evaluated using the current window location, then pulls the values into their least-mean squares values for the current window. Of course, as with any tracking method, if the current window location is too far from the previous one, any one track may jump from one scatterer to another or be lost entirely.

To take care of the problem that one scatterer may become so weak that the gradient algorithm pulls its track into coincidence with that of another, a clean-up routine is necessary to remove duplicates. If the duplicates are not removed, the subsequent operation to evaluate the scatterer amplitudes a_s faces the problem of an ill-conditioned matrix. The removal is readily carried out by retaining only one of any group of tracks for which the effective coordinates (λ_s, ν_s) are mutually closer than a given minimum, using a 2-norm measure of proximity.

Another problem to be addressed is that of accommodating new scatterers as their amplitudes rise into significance. This is handled by constructing the lists of initial estimates of the τ_s and ν_s presented to the gradient algorithm in the current window location by simply appending to the lists evaluated in the previous window location the lists of all the significant scatterers appearing in the new window location, whether old or new. Most of the scatterers will then be represented twice in the lists, but after they have been subjected to the gradient algorithm, those represented twice will end up having essentially the same effective coordinates and the new ones are then eliminated by the clean-up routine. A unique identification number is attached to every scatterer found in each new data window location. The procedure described ensures that as long as the scatterer remains significant, it will retain this unique ID. And, of course, by virtue of the duplication of the lists, most IDs will end up being eliminated.

TOTAL LEAST SQUARES PROCESSING

One more potential noise-reduction technique that should be mentioned is that of so-called total least squares (TLS) processing [2,12,13]. The idea is that all the columns of all three of the mutually shifted matrices U_0 , U_i , U_f are closely related. Specifically, each of these columns, apart from the additive noise it contains, is a different linear combination of the same S column vectors, one for each of the S scatterers. Accordingly, cleaner versions of them can be evaluated by using singular value decomposition to factor the block matrix $[U_0 \ U_i \ U_f]$ into the form

$$[U_0 \quad U_t \quad U_f] = U \Sigma \begin{bmatrix} V_{11}' & V_{21}' & V_{31}' \\ V_{12}' & V_{22}' & V_{32}' \\ V_{13}' & V_{23}' & V_{33}' \end{bmatrix}$$

and then using the trio of square matrices V_{11}' , V_{21}' , and V_{31}' in place of the trio U_0 , U_t , U_f in (3) to evaluate the matrices R_t and R_f .

In tests of radar imaging conducted so far, the inclusion of TLS processing has produced no solid improvement. This is probably attributable to the fact that column-relatedness is already a key property in the processing, and is therefore essentially fully exploited without TLS processing.

REFERENCES

1. R. Roy and T. Kailath, "ESPRIT — Estimation of Signal Parameters via Rotational Invariance Techniques," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 37, no. 7, pp. 984-995, July 1989.
2. B. Ottersten, M. Viberg, and T. Kailath, "Performance Analysis of the Total Least Squares ESPRIT Algorithm," *IEEE Trans. Signal Processing*, vol. 39, no. 5, pp. 1122-1135, May 1991.
3. J. Choi, N. Wang, L. Peters, Jr., and P. Levy, "Near Axial Backscatterering from Finite Cones," *IEEE Trans. Antennas Propagat.*, vol. 38, no. 8, pp. 1264-1272, Aug. 1990.
4. J. B. Keller, "Backscattering from a Finite Cone," *IEEE Trans. Antennas Propagat.*, vol. AP-8, no. 2, pp. 175-182, March 1960.
5. S. Y. Kung, K. S. Arun, and D. V. Bhaskar Rao, "State-space and Singular-value Decomposition-based Approximation Methods for the Harmonic Retrieval Problem," *J. Opt. Soc. Am.*, vol. 73, no. 12, pp. 1799-1811, Dec. 1983.
6. J. E. Piou, J. T. Mayhan, K. M. Cuomo and M. L. Burrows, "Coupled Multi-input Multi-output (CMIMO) Systems: 2-D Spectral Estimation Techniques" (submitted for publication)
7. J. Mayhan, M. L. Burrows, K. M. Cuomo, and J. E. Piou, "High Resolution 3D "Snapshot," ISAR Imaging and Feature Extraction" *IEEE Trans. Aero. and Electron. Systems*, vol. 20, no. 2, pp. 630-642, April, 2001.
8. J.G. Proakis, C. M. Rader, F. Ling, and C. L. Nikias, *Advanced Digital Signal Processing*, Macmillan, 1992, pp. 535-538.
9. S. L. Borison, S. Bowling, and K. M. Cuomo, "Superresolution Methods for Wideband Radar," *Lincoln Laboratory J.*, vol. 5, no. 3, pp. 441-461, 1992.
10. G. R. Benitz, "High-Definition Vector Imaging for Synthetic Aperture Radar," *Conference Record of the Thirty-First Asilomar Conference on Signals, Systems & Computers*, vol. 2, M. P. Fargues and R. D. Hippenstiel, eds., IEEE Computer Society, Los Alamitos, CA, 1998, pp. 1204-1208.
11. F. Vanpoucke, M. Moonen, and Y. Berthoumieu, "An Efficient Sub-Space Algorithm for 2-D Harmonic Retrieval," *Proceedings of ICASSP*, vol. 4, pp. 461-464, IEEE, New York, 1994.
12. G. H. Golub and C. F. Van Loan, *Matrix Computations*, Johns Hopkins, 1996, p. 395ff.
13. P. Strobach, "Total Least Squares Phased Averaging and 3-D ESPRIT for Joint Azimuth-Elevation-Carrier Estimation," *IEEE Trans. Signal Processing*, vol. 49, no. 1, pp. 54-62, Jan. 2001.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 2 August 2002		3. REPORT TYPE AND DATES COVERED Technical Report
4. TITLE AND SUBTITLE Two-Dimensional ESPRIT with Tracking for Radar Imaging and Feature Extraction			5. FUNDING NUMBERS C—F19628-00-C-0002	
6. AUTHOR(S) M.L. Burrows				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Lincoln Laboratory, MIT 244 Wood Street Lexington, MA 02420-9108			8. PERFORMING ORGANIZATION REPORT NUMBER TR-1084	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Space & Missile Defense Command P.O. Box 1500 Huntsville, AL 35807-3801			10. SPONSORING/MONITORING AGENCY REPORT NUMBER ESC-TR-2001-075	
11. SUPPLEMENTARY NOTES None				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) <p>ESPRIT processing appears to be the best of the known spectral-analysis techniques. It provides the highest resolution and has no spectral splatter. By applying matrix eigenstructure analysis, it gives a direct answer to the direct question "What frequencies, real or complex, are present in the data and what are their amplitudes?" Conventional Fourier techniques, as well as some of the other higher-resolution methods, answer the less direct question "What amplitudes, applied to a set of regularly-spaced real frequencies, best represent the data?" Then comes the problem of interpreting the amplitudes.</p> <p>These attributes of ESPRIT, in the two-dimensional version described here, make it a natural for radar signal processing, where it answers the need for high-resolution imaging, free of sidelobes in range and Doppler, and for high-fidelity target feature extraction. For example, the uncertainty in the scatterer-center locations in an ESPRIT image extracted from high-quality static-range radar data collected over a bandwidth of 1 GHz is just a few millimeters; for conventional Fourier processing of the same data the uncertainty is many centimeters. The signature of the base edge of a perfectly conducting cone extracted from static-range data by ESPRIT agrees accurately with the signature predicted by edge-diffraction theory.</p> <p>This report starts with a mathematical model for the radar data, describes a technique for "resampling" the data to achieve a more perfect fit with the ESPRIT data model, summarizes the two-dimensional ESPRIT algorithm itself, and presents examples of its performance. The appendix covers the details of this least-mean-square version of ESPRIT, including an enhancement that allows the scatterers to be tracked individually. The algorithm properly distinguishes between target locations having one coordinate in common, and it automatically associates correctly in pairs each entry in the list of ranges with the corresponding entry in the list of range rates.</p>				
14. SUBJECT TERMS radar imaging high-resolution ESPRIT spectral analysis			15. NUMBER OF PAGES 32	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT Same as Report	